IFW

520.43077X00

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicants: Shuji NAKAMURA et al.

Serial No.: 10/648,291

Filed: August 27, 2003

For: STORAGE SYSTEM

PETITION TO MAKE SPECIAL
UNDER 37 CFR 1.102(d) and MPEP. §708.02, VIII

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

December 29, 2004

Sir:

1. Petition

Applicants hereby petition to make this application **Special**, in accordance with 37 CFR §1.102(d) and MPEP 708.02, VIII. The present invention is a new application filed in the United States Patent and Trademark Office on August 27, 2003 and as such has not received any examination by the Examiner.

2. Claims

Applicants hereby represent that all the claims in the present application are directed to a single invention. If upon examination it is determined that all the claims presented are not directed to a single invention, Applicants will make an election without traverse as a prerequisite to the granting of special status.

3. Search

Applicants hereby submit that a pre-examination search has been made by a professional searcher, (a copy of which is attached), in the following classes and subclasses:

Class	Subclass
711	111, 113, 119, 162
714	5, 6

4. Copy of References

A listing of all references found by the professional searcher is provided on a Form PTO-1449 and copies of the references and the Form PTO-1449 are submitted as part of an Information Disclosure Statement (IDS) filed on even date.

5. Detailed Discussion of the References and Distinctions Between the References and the Claims

Below is a discussion of the references uncovered by the search and cited in the IDS filed on even date that appear to be most closely related to the subject matter encompassed by the claims of the present application, and which discussion particularly points out how Applicants' claimed subject matter is distinguishable over those references. All other references uncovered by the search and cited in the IDS filed on even date are **not** treated in detail herein.

a. Detailed Discussion of the References

U.S. Patent No. 5,771,367 (Beardsley) discloses a storage controller and method for improved failure recovery using cross-coupled cache memories and nonvolatile stores. The controller includes a first cluster for directing data from a host computer to a storage device and a second cluster for directing data from a host computer to a storage device. A first cache memory is connected to the first cluster and a second cache memory is connected to the second cluster. A first nonvolatile memory is connected to the second cluster and a second nonvolatile memory is connected to the first cluster. Data is directed to the first cache and backed up to the first nonvolatile memory. In the event of failure of the first nonvolatile memory, data is recovered from the first cache memory. In a particular embodiment, in the event of failure of the first nonvolatile memory, data is directed from the first cache memory to the second cache memory (see, abstract, figures 1-7, column 3, lines 20-45, columns 4 and 5).

U.S. Patent No. 5,809,543 (Byers) discloses a fault tolerant extended processing complex for redundant nonvolatile file caching. In the system, to provide redundancy and resiliency against error, the file data signals to be cached are stored in duplicate files in separate portions of the redundant nonvolatile file cache storage. All of the control and access circuitry is duplicated such that a complete set is applicable to each half of the redundant nonvolatile file cache storage. By providing the redundant copy of the cached file signals, a failure in any of the control or access circuitry will not result in system failure, but will accommodate recovery of the cached file data signals from the portion of the redundant nonvolatile file cache storage to which access has

been maintained (see, abstract, figures 1-37, column 6, lines 45-67, column 7, lines 15-55, and column 8, lines 1-55).

U.S. Patent No. 6,006,342 (Beardsley) discloses a failover and failback system for handling failures in a storage controller interfacing between a plurality of host systems and direct access storage devices. The storage controller directs data from the host systems through first and second data paths in the storage controller to a DASD. A first processor, first non-volatile memory unit (NVS), and a first cache are associated with the first data path, and a secondary processor, a second NVS, and a second cache are associated with the second data path. A bridge provides communication between the first processor and the second NVS and the second processor and the first NVS.

U.S. Patent No. 6,279,078 (Sicola) discloses an apparatus and method for synchronizing a cache mode in a dual controller, dual cache memory system operating in a plurality of cache modes. The cache memory system has two controllers and two cache modules and operates in a non-mirror cache mode and a mirror cache mode. Data indicating the cache mode to be used is stored as metadata in the cache modules. The metadata in the cache modules is detected to determine the cache mode. Lock signals in one of the controllers are set in accordance with the cache mode to set the cache mode state of the controller. With the cache mode state being a mirror or non-mirror state, the other controller copies the lock state from the first controller to synchronize both controllers in the same cache mode state. After a failure of the second controller, the first controller may lock access to both caches to recover data

previously accessed by the second controller. The second controller restarts and copies the cache mode state from the first controller, so that both controllers return to the cache mode state prior to the failure of the second controller (see, abstract, figures 1-3, column 1, liens 57-67, column 2, lines 1032, 50-67 and column 3).

U.S. Patent No. 6,760,765 (Asai) discloses a cluster server device with which data can be distributed to terminal equipment while performing the optimal distribution of load to plural cache servers even when a fault occurs in any cache server. A cluster control part distributes a request from terminal equipment corresponding to the loads of cache servers. When requested data (streaming data) exists in a streaming data storage part, the relevant data are distributed to the terminal equipment by the cache server, and when such data do not exist, data distributed from a contents sever are distributed to the terminal equipment.

U.S. Patent Publication No. 2002/0133735 (McKean) discloses a system and method for efficient failover/failback techniques for a fault-tolerant data storage system. The system provides a method for an efficient failback technique in a data storage system utilizing a dual-active controller configuration for minimizing a delay in responding to input/output (I/O) requests from a host system following a controller failure. A stripe lock data structure is defined to maintain reservation status or stripe locks of cache lines within data extents that are part of a logical unit or storage volume. When a controller fails, dirty cache line data of a failed controller is taken over by a survivor controller. The stripe lock data structure is used to process I/O requests from a host system, by the failed controller. The data storage system functions in a single-

active configuration until the dirty cache line data is flushed to one or more storage volumes by the survivor controller. Once the dirty cache line data is flushed to the system drive, the data storage system continues processing the host I/O requests in the dual-active or multiply-active configuration. Hence, McKean's system discloses the teaching of dual control units each having a memory for implementing a fault-tolerant data storage system.

U.S. Patent Publication No. 2004/0010659 (Inoue) discloses an external storage subsystem capable of implementing an error-tolerable, highly-reliable cache function. A plurality of mutually independent cache units 80 and 81 and non-volatile memory units 90 and 91 are provided in a disk control unit 2 between a central processing unit 1 and a magnetic disk unit 3. A plurality of channel units 60 and 61 that control the transfer of data to or from the central processing unit 1 and a plurality of control units 70 and 71 that control the transfer of data to or from the magnetic disk unit 3 are independently connected to the plurality of cache units 80 and 81 and non-volatile memory units 90 and 91 via data buses and access lines.

U.S. Patent Publication No. 2004/0153727 (Hicken) discloses a method for recovering redundant cache data of a failed controller and reestablishing redundancy. In a cache-redundant data storage system having at least two storage controllers, with each of the storage controllers comprising a primary cache memory and a second cache memory, Hicken teaches a method for recovering cache data of a failed redundant storage controller and reestablishing redundancy. The method includes arranging at least two storage controllers into pairs, such that the cache data of the

primary cache memory of each of the storage controllers in the pair is mirrored in the secondary cache memory of the other storage controller in the pair. Upon a failure of a storage controller in a pair, the failure is detected, and, in response, a structured list of cache tags is created in the non-failed storage controller. The primary cache memory is and the secondary cache memory of the non-failed storage controller are flushed, and an available secondary cache memory is configured to function as a redundant cache memory for the primary cache memory of the non-failed storage controller, such that cache data in the primary cache memory of the non-failed storage controller is mirrored in the available secondary cache memory.

b. Distinctions Between the References and the Claims

The present invention as recited in the claims filed are not taught or suggested by any of the above noted references whether taken individually or in combination with each other or in combination with any of the other references now of record.

The present invention as recited in the claims is directed to a storage system connected to a computer that includes: a first control unit, a second control unit, a third control unit and plural storage units, wherein the first control unit, the second control unit and the third control unit each has a memory, and the first control unit stores data received from the computer in the memory possessed by the first control unit and the memory possessed by the second control unit.

The above described features of the present invention, particularly a storage system that includes: a first control unit, a second control unit, a third control unit and plural storage units, where the first control unit, the second control unit and the

third control unit each has a memory, and the first control unit stores data received from the computer in the memory possessed by the first control unit and the memory possessed by the second control unit, are not taught or suggested by any of the references of record whether taken individually or in combination with each other.

6. Fee (37 C.F.R. 1.17(i))

The fee required by 37 C.F.R. § 1.17(i) is to be paid by:

[X] the Credit Card Payment Form (attached) for $130.00.

[ ] charging Account _____ the sum of $130.00.

A duplicate of this petition is attached.

Please charge any shortage in fees due in connection with the filing of this paper, including extension of time fees, or credit any overpayment of fees, to the deposit account of Mattingly, Stanger & Malur, P.C., Deposit Account No. 50-1417 (520.43077X00).

Respectfully submitted,

MATTINGLY, STANGER & MALUR, P.C.

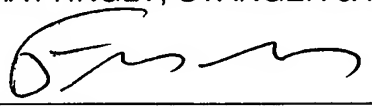

Frederick D. Bailey
Registration No. 42,282

FDB/sdb
Enclosures